Security Issues in All-Optical Networks

Muriel Médard, Douglas Marquis, Richard A. Barry, and Steven G. Finn
Massachusetts Institute of Technology Lincoln Laboratory

Abstract

All-optical networks are emerging as a promising technology for terabit per second class communications. However, they are intrinsically different from electro-optical networks, particularly because they do not regenerate signals in the network. The characteristics of all-optical network components and architectures manifest new and still unstudied security vulnerabilities but also offer a new array of possible countermeasures.

An all-optical network (AON) is a network where the user-network interface is optical and the data does not undergo optical-to-electrical conversion within the network. AONs are attractive because they promise very high rates, flexible switching, and broad application support. Currently, optical transmission links supporting 30–40 Gb/s are commercially available, 100 Gb/s products have been announced, and terabit-per-second AONs have been demonstrated in the laboratory [1, 2]. AONs offer the potential to tap economically this large capacity because fiber optic transmission technology is progressing faster than electronic switching technology, and because optical switching technology is maturing to the point where it may become the economical choice in certain applications. Research testbeds [3–7] have demonstrated basic AON functionality with transmission rates of over 100 Gb/s in local and metropolitan area networks. Figure 1 shows where AONs may fit in a communications infrastructure.

The emergence of these networks coincides with a burgeoning use of networked information for education, commerce, health care, national defense, and many other endeavors that promise continued growth for decades. The transition from electro-optic networks to all-optical networking in some sectors thus may come at a time when societal reliance on networks is greater than its already significant level. Assured access to these networks, in a private and reliable manner and with appropriate service guarantees, is clearly very important and has motivated our study of the security of AONs. The study of security for AONs is fairly new, but already presents many interesting issues. Note that many of the security problems present in traditional electronic or electro-optic networks are still to be found in AONs. Moreover, many of the traditional security mechanisms, such as end-to-end encryption, may be applicable to AONs and, in certain cases, may be even more necessary.

In this article, we concentrate on the *physical* security of data in AONs. Physical security ensures that the data have a minimum privacy and quality of service (QoS), and that users are informed when such privacy or QoS are unavailable or when failure to deliver the required privacy or QoS has occurred. This is in contrast to *semantic* security, which protects the meaning of the data even when an attacker has access to the received stream. Semantic security thus belongs to the realm of cryptography, which is outside the scope of this article.[1] We concentrate on physical security because it directly impacts the construction of AONs, and physical security in AONs differs substantially from that in electro-optical networks. Note that, although we separate physical and semantic security, some techniques (e.g., spread-spectrum) may address both problems. Note also that cryptography cannot protect against physical-layer attacks which attempt to deny or degrade service.

We focus on two types of attacks on the physical security of the network: service disruption (SD), which prevents communication or degrades QoS, and tapping, which compromises privacy by providing unauthorized users access to data which may be used for eavesdropping or traffic analysis. We will see that both SD and tapping may be achieved by gaining access to the network through authorized or unauthorized entry points. An unauthorized entry point may be obtained by cutting a fiber, for example. An authorized entry point may allow an attacker to enter the network under the guise of a legitimate user by sending a nefarious signal instead of or in addition to a legitimate communication. An important attack on physical security not addressed is a software attack on the network management and control system. Although AONs require a network management and control system, we do not consider the problem of compromising these systems, for such

[1] *Semantic security schemes which rely specifically on optics, such as quantum cryptography or homodyne scrambling, exist in research environments (see [8, 9]).*

0890-8044/97/$10.00 © 1997 IEEE

a software attack is not conceptually different in AONs from electro-optic networks. Finally, we focus on issues of security against attacks, not reliability under normal operation.

The second section is an overview of AONs, describes some of their fundamental components, and highlights differences between AONs and traditional electro-optic networks which affect physical security. The third section describes some of the vulnerabilities particular to AONs and presents attacks which perform SD and tapping by exploiting these vulnerabilities. The fourth section discusses security countermeasures in AONs. Our conclusions in the final section summarize the main issues.

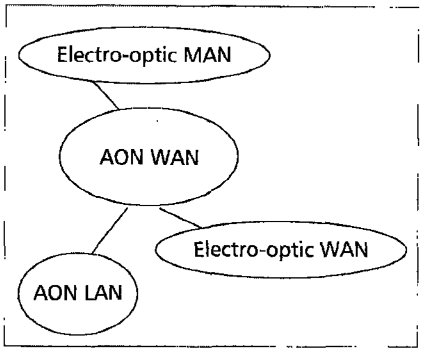

■ Figure 1. *Role of an AON in a communications infrastructure.*

All Optical Networks

General Description

AONs exist in today's research environments (and in early commercial systems) in two types: time-division multiplexed (TDM) networks (generally employing 100+ Gb/s soliton transmission, possibly in multiple bands), and wavelength-division multiplexed (WDM) networks, similar to the frequency multiplexing used in some radio frequency (RF) transmission systems. Our discussion will concentrate on AONs employing WDM. WDM networks divide the fiber bandwidth into optical *wavebands*, each of which carries information at rates up to ~10 Gb/s. These WDM waveband channels are routed through the network using wavelength- or frequency-selective devices and optical switches. WDM transmission systems are currently the most promising for optical networks because WDM technology is more mature than optical TDM [10].

To establish an optical circuit between terminals, a circuit is assigned a waveband which is routed through the AON. This process, called *wavelength routing*, is illustrated by a simple example shown in Fig. 2. The user terminal in the upper left corner may transmit over wavelengths λ_1–λ_8. A wavelength router separates λ_1–λ_5 from λ_6–λ_8. In the upper right corner, a wavelength-selective switch (WSS) sends λ_1 to one user destination, λ_2 and λ_3 to another, while λ_4 and λ_5 continue on in the network. Similarly, the bottom left add-drop multiplexer (ADM) passes λ_6 to another network node, and sends λ_7 and λ_8 to a user. Not shown in Fig. 2 are possible wavelength changers, which may allow a circuit to terminate on a different waveband than that on which it originated. For simplicity of exposition, we do not consider wavelength changers in this article nor multiplexing schemes such as packetization, time-division multiple access (TDMA), or code-division multiple access (CDMA), which may be used to share a waveband within the network. Rather, we focus on the situation, analogous to the use of AONs for backbone traffic, where circuits occupy an entire waveband.

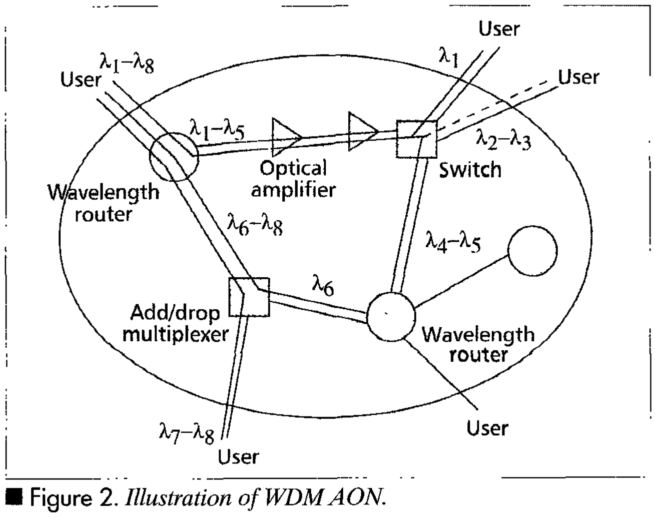

■ Figure 2. *Illustration of WDM AON.*

Although we have chosen a small mesh configuration to illustrate wavelength routing, it is important to keep in mind that many different optical network architectures exist. For instance, the ATT/Digital Equipment Coporation (DEC)/Massachusetts Institute of Technology (MIT) AON is built on a hierarchical architecture which supports both full waveband circuits and time sharing of wavebands [3]. The Optical Network Technology Consortium (ONTC) testbed is composed of two interconnected rings and uses an OC-3c synchronous optical network (SONET) format [6]. The Research and Development in Advanced Communication Technologies in Europe (RACE) Multiwavelength Transport Network (MTWM) network is a mesh network which carries asynchronous transfer mode (ATM) as well as dark fiber traffic [7]. The Multiwavelength Optical Network (MONET) project includes ring, cross-connect mesh, and star networks [5]. Despite the variety of architecture and service implementations, all of the AON testbeds share common component classes, in particular fibers, optical amplifiers, and wavelength-selective nodes (routers or switches). Thus, many features of physical security are common to all AONs and may be discussed independent of the specific implementations.

Key Differences between AONs and Traditional Networks

Some AON features can be used by an attacker to analyze traffic, eavesdrop, degrade QoS, or deny service altogether. The following is a list of features and vulnerabilities that distinguish the security characteristics·of AONs from electronic and electro-optic communication networks.

The very high data rates enabled by all-optical technology have two important security ramifications. First, even attacks that are short and infrequent can result in large amounts of data being corrupted or compromised. Second, end users may choose to retain protocols designed for slower electronic networks (e.g., routed TCP/IP — Transmission Control Protocol/Internet Protocol — over telephone company infrastructure). While such protocols perform well in the domains for which they were intended, the use of such protocols at link, network, and transport layers at very high bit rates over large distances can allow effective service denial attacks using sporadic or relatively low-power methods. Such attacks may well be very difficult to detect.

Wavelength routing provides a *transparent* circuit between terminals. Signals that are routed transparently are amplified but not regenerated at any network component. This transparency is desirable in many respects; for instance, it provides a simple way for heterogeneous terminals to share network resources. Certain wavebands can carry analog signals, while other wavebands are simultaneously used for digital signals. Thus, different terminals may use different modulation formats, and terminals may be upgraded without wholesale network reconfiguration. However, transparency raises many security vulnerabilities which do not exist in electro-optic networks. Malicious signals may be designed to pass through transparent components, disabling portions of a network to which they would not have had access in a regenerated system. In a network with regeneration, an intermediary node will not propagate an anomalous signal. Rather, a regenerating node would generally discard an anomalous signal or may misinterpret one and generate an erroneous signal. However, the output of the regenerating node will comply with a certain set of signaling schemes and protocols. Thus, an ATM switch may be jammed to cause certain cells to be dropped or erroneously transmitted. However, its output will continue to be ATM cells transmitted with a certain modulation, signal intensity, within a certain bandwidth, and so on. As a rule, regeneration cordons off certain types of attacks.

The inability of AONs to reconstruct data streams at nodes and repeaters within transparent networks also significantly complicates segment-by-segment testing of communication links, since the transparent network node by design will not understand the modulation and coding of the signal. This makes attack and fault localization difficult using current segment-by-segment methods. The combination of transparency and multiple WDM streams per fiber gives would-be jammers an ability to foul multiple streams of traffic by rapidly switching jammer power among streams.

Although rapidly maturing through heavy commercial and governmental investment, many AON components are less mature and robust than their electronic or electro-optic network analogs. Relatively high crosstalk between WDM channels within existing components appears particularly important to the security problem. Crosstalk can be exploited either to tap communications or to perform SD by injecting malicious signals into a network. Nonlinearities in fibers and devices can lead to undesirable cross-modulations which may cause SD or subtle tapping attacks. Optical amplifiers under attack by jammers may cease to amplify and thus lead to SD. Although many AON components such as fibers and amplifiers also exist in electro-optic networks, transparency renders their vulnerabilities more important for physical security than in a network with regeneration.

The unmasking of these vulnerabilities could be interpreted as an intrinsic weakness of AONs, especially to SD attacks; such an interpretation would be premature. Countermeasures to attacks attempting to exploit nearly all of the above features are available, although those countermeasures are not always intuitive or traditional. The above list simply points out that the network security framework for AONs is different from that for electronic or electro-optic networks. The problem is new, and requires new consideration and new solutions.

Vulnerabilities and Attacks

In this section we discuss particular vulnerabilities of certain key components and subsystems of AONs, and how attackers can make use of these vulnerabilities to perform SD or tapping. All of the attacks discussed below only require physical access to a fiber of the network within a reasonable distance of a node. While the distance will depend heavily on the form of the attack and the specific type of fiber used, a distance of 10–20 km might be a typical range. Some of the attacks can actually be performed by a user at a legitimate network access point.

We assume that it is unreasonable to expect that metropolitan area network (MAN) or wide area network (WAN) fibers will be completely physically secure; thus, we believe that for the foreseeable future access to fibers will generally be available to the resourceful. We also assume a *mixed traffic environment* where trusted signals, (i.e., signals carrying sensitive, possibly classified, information) may share links and/or nodes with calls from untrusted signals (i.e., signals launched by potential attackers). The most secure solution would be to separate trusted and untrusted users on different fibers using different nodes, thereby essentially providing two different networks. However, there are a number of arguments against this approach. First, for an arbitrary topology this may not always be possible. Second, even if the initial topology allowed separation of users, link failures may force mixed traffic. Third, the security level of a link, node, or user may change dynamically in response to anomalous behavior. Fourth, there may actually be multiple levels of security, and as the number of security levels increases, the practicality of physically separate networks diminishes.

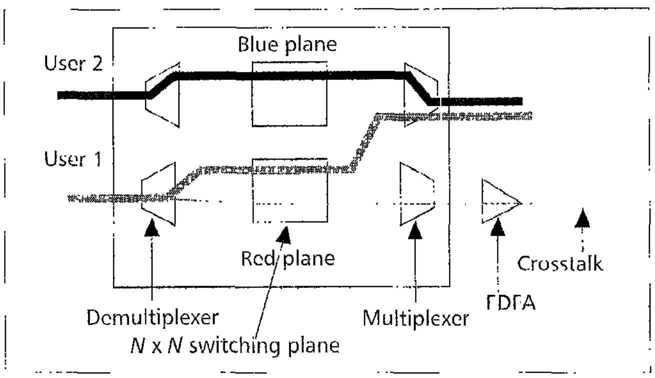

■ Figure 3. *WSS with crosstalk.*

Service Disruption Attacks

We discuss SD attacks in the context of three main AON components: fibers, amplifiers, and wavelength selective switches (WSSs). Although, as discussed previously, implementations vary among AONs, these classes of devices are fundamental building blocks.

Fibers — Fibers ideally propagate light on different wavelengths with only frequency-dependent delay (dispersion) and attenuation. Fibers have very low loss (0.2 dB/km) and are fairly linear under normal operating conditions.

Under normal operating conditions, there is a negligible radiation of power from the fiber. This is in contrast to other waveguide media, such as coaxial cable. However, like coax, unshielded fiber presents little security against an attacker with physical access; service is easily disrupted by cutting or in some way disrupting the fiber. A less disruptive attack is to bend the fiber slightly so that light may be radiated into or out of the fiber [11]. By using this or related techniques, service can be interrupted on a single wavelength by injecting light at that wavelength (in-band jamming) without breaking or otherwise disrupting the fiber. This attack is difficult to localize.

Under high power input or long distances, fibers do exhibit certain nonlinear characteristics that cause signals on different wavelengths to interact. For instance, a signal on one wavelength may cause amplification or attenuation of a signal on another wavelength through cross-phase modulation and Raman gain effects [12]. Thus, there are crosstalk effects among wavelengths in a fiber which may be exploited by a sophisticated attacker.

Optical Amplifiers — Rare-earth doped optical amplifiers (OAs), especially erbium doped fiber amplifiers (EDFAs), have revolutionized fiber optic communication and are a critical and necessary component for essentially all modern fiber optic communication systems. OAs transparently amplify optical signals and are used to restore the optical signal power to an acceptable level. The most common usage of OAs is on fiber optic links, where they are used to compensate for fiber attenuation. OAs are also used in network nodes to compensate for losses from optical devices — there are many EDFAs in most AONs.

The nature of EDFA operation in WDM communication links and nodes can lead to a phenomenon known as *gain competition*, whereby multiple independent WDM wavebands share a limited pool of available upper-state photons within the fiber [13]. The result is that a stronger user (possibly an attacker) can deprive a weaker user of photons, thus reducing the weaker user's gain. This gain competition, combined with the fact that fiber has extremely low loss, means that EDFAs are susceptible to power jamming from remote locations. The severity of this attack depends on the distance of the attacker from the OA, the OA device specifications, and the network architecture. In some instances, it may be possible for a malicious user to deny service to many other users from a legitimate network access point.

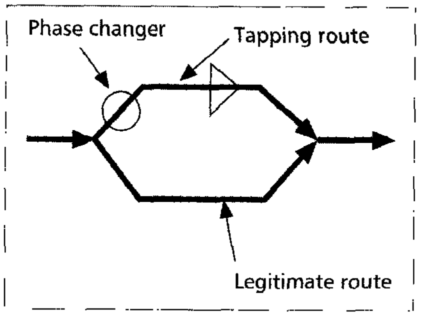

■ Figure 4. *An example of correlated jamming.*

WSSs — WSSs route signals of different wavelengths to different outputs. Two building blocks of WSSs are wavelength demultiplexers and multiplexers: a demultiplexer separates signals of different wavelengths, each onto its own fiber; a multiplexer performs the opposite function. Multiplexers are relatively inexpensive and are commercially deployed on a small scale. Figure 3 shows a typical WSS. The WSS illustrated has demultiplexers at the input, which send all the signals of a particular wavelength to one of two switching planes. The switching planes permute the signals, which are then multiplexed on the output fibers. To make up for device losses, OAs may be used on the inputs, outputs, or both. OAs may also be used in the switching planes. In fact, some switch architectures use OAs as optical crosspoints. Other devices for routing and switching include star couplers, wavelength filters, ADMs (for both adding/dropping fibers or for adding/dropping wavelengths), wavelength changers, and wavelength routers.

Current state-of-the-art examples of WSSs have significant crosstalk levels. Crosstalk causes signals to leak onto unintended outputs and permits inputs to cause interference on other optical signals going through these devices. The level of crosstalk greatly depends on the particular components and architecture of a WSS. Figure 3 shows the crosstalk associated with a WSS going from user 1 through the WSS and out the bottom output port. State-of-the-art multiplexers and demultiplexers may have crosstalk levels of –20 dB (e.g., acousto-optical switch) to –30 dB (most other switch types). Note, however, that crosstalk is additive, and thus the aggregate effect of crosstalk over a whole AON may be more nefarious than a single point of crosstalk.

Although –30 dB crosstalk is adequate for most communication purposes, it permits several types of attacks. An attacker could use this small level of crosstalk to disrupt or deny service. For service denial, the attacker injects a very strong signal into the device (the attacker may be physically distant from the node by using a network fiber to inject the signal). Although only a small fraction of it leaks onto another channel, a sufficiently powerful signal modulated in a malicious way can be highly disruptive. A possible attack on an on-off keyed (OOK) signal would occur if the jammer chose a bit rate which was sufficiently fast that the receiver's threshold could not adjust, but not so fast as to be filtered out by the receiver filter. In another possible attack, an EDFA in a WSS, such as in Fig. 3, could be jammed if user 1 uses crosstalk to insert a jamming signal onto the bottom output fiber, thus jamming any users whose signals share that EDFA. The effects of jamming an OA in a switch can be severe. A sophisticated attacker could potentially disrupt an *entire* network node remotely with physical access to a single input fiber of that node. This is a very different situation than in electro-optic networks, where regeneration would isolate the attack. The attacker accomplishes an attack by injecting light of multiple wavelengths, or by rapidly frequency hopping from one wavelength to another for the purpose of jamming amplifiers on as many output ports as possible. The extent of disruption depends on the location of the attacker, the switch architecture, and the state of the switch (i.e., which wavelengths are routed to which outputs). A particularly nefarious enemy with cooperation from malicious users may attempt to put the switch in a vulnerable state by having his cohorts make seemingly legitimate service requests.

Tapping Attacks

Fibers and EDFAs — An attacker with physical access to the fiber can retrieve part of a signal with little disruption by slightly bending the fiber. Also, at high signal levels (e.g., at the output of an OA), fibers exhibit some crosstalk which may be used for tapping by copropagating a signal on the fiber.

In some OAs, the gain competition among signals occurs very rapidly (i.e., at the modulation rate). In these scenarios, tapping can be achieved by observing cross-modulation effects.

Note that tapping can be combined with jamming for a particularly effective SD attack. Since delays vary extremely slowly with respect to data rates in AONs, an attacker can tap a signal and then inject a signal downstream of the tapping point. Such an attack, which we call *correlated jamming*, is particularly pernicious when the signal-to-noise ratio (SNR) of the attacked user is relatively low. Note also that the effect of a correlated jamming attack is greater than the effect of an uncorrelated jamming attack for the same jammer power. The attacker can subtract the signal partially and replace it with noise. He may also introduce multiple different delays to introduce repeat-back jamming, which may appear at the receiver as naturally occurring multipath. The correlated jammer attack is illustrated in Fig. 4.

WSSs — Owing to crosstalk, WSSs are vulnerable to tapping. Consider in Fig. 3 that user 2 is tapped by an attacker with access to the bottom output fiber.

A simple and seemingly effective software solution to this problem is to amplify only signals which are "supposed" to be there (i.e., registered with the network management system). This requires more switches or OAs to independently amplify different wavelengths. In Fig. 3, the crosstalk would not be amplified. However, an attacker can thwart this measure by making his own request. This is illustrated in Fig. 5 by user 1's call. The attacker is user 1. The attacker's call, which is also on Blue, is routed onto the bottom fiber by the WSS. This means that the call from user 1 is superimposed on the leakage of the call from user 2. Since there is now a call on the bottom fiber on Blue, the network management software will

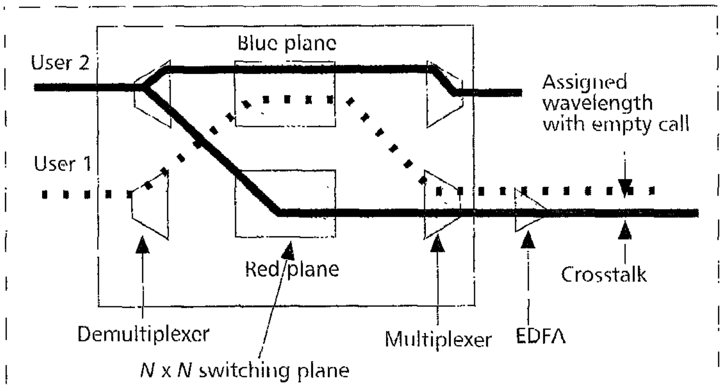

■ Figure 5. *An example of a tapping attack using WSS.*

not turn off the amplifier. In order to compensate for unexpected or unpredictable loss, the network control algorithms may furthermore provide extra amplification to the weak signal in response to a small measured signal power. However, if user 1 does not transmit any power on Blue, the only signal received from user 1's call is the amplified leakage from user 2's call, which may be sensitive. The network control algorithms unwittingly amplify and route the call into the hands of an attacker. There are countermeasures to this attack, which in themselves may have more subtle vulnerabilities.

These examples demonstrate the variety of possible attacks, the interplay between countermeasures for different attacks, and the need for a comprehensive set of countermeasures controlled with an appropriate network management system.

Countermeasures

The philosophy we have adopted in examining network security schemes has three components: prevention, detection, and reaction. We assume the existence of a network management and control entity and of a security status database, which may be centralized or distributed.

While some of the mechanisms are applicable to different types of networks, many issues are peculiar to AONs. The large data rates in AONs entail that, even if the network was under attack for a few seconds, terabits per second of information may be compromised. Thus, detection and reaction should be rapid. Since the AON does not perform signal regeneration, it cannot make use within the network of such means of error detection as CRCs. Thus, AONs must be able to detect covert SD attacks, which disrupt the signal enough to reduce QoS without precluding all communication, as well as overt attacks, such as altogether eliminating the signal or overwhelming it by strong, protracted jamming. AONs must also be able to detect tapping attacks that negligibly affect the received SNR as well as those that steal a significant portion of the signal.

Prevention

Prevention methods available may be broadly divided into three categories: techniques that reduce the vulnerabilities which are intrinsic to the hardware, transmission schemes that are inured against certain attacks, and protocols and architecture designs adapted to AONs.

We first consider prevention of SD through hardware measures. The device requirements for communications in the absence of attacks may be different than those for security. For instance, a certain level of crosstalk may be tolerated when all users remain within certain power limits, but may become unacceptable if one user sends a strong jamming signal or performs correlated jamming. Some hardware measures seek to mitigate jamming by ensuring that the signals in the network comply with certain specifications. An optical limiting amplifier (OLA), which amplifies but limits the output power to a certain maximum, may be employed to enforce certain limits on power. Such limits will hinge on the security requirements and the levels of crosstalk. Bandlimiting filters may be used to discard signals outside a certain bandwidth, thus preventing out-of-band gain competition attacks in OAs.

Hardware measures also can harden against tapping. Preventing a physical tap into the fiber can be achieved by physical strengthening, alarming the cladding, or attempting to detect small losses of power due to tapping. All of these measures except for the detection of loss of power require substantial changes in the existing infrastructure and would entail significant expense [14]. Although alarming the cladding, using fiber that breaks when bent, or strongly shielding the fiber are viable options for a local area network (LAN) within a secure enclave, they will not in general be possible in a MAN or WAN. Moreover, securing the fiber against physical tapping does not protect against tapping via crosstalk. Devices with lower crosstalk will mitigate both SD and tapping attacks.

Hardware measures must also be considered in the context of the component impact, layout of the components, and cost trade-offs. For instance, introducing OLAs may allow us to withstand higher crosstalk levels in WSSs; therefore, the cost of introducing OLAs must be compared against the cost of improving the crosstalk in WSSs. Another trade-off arises if untrusted users can obtain a tapping channel, by means of crosstalk, which is equivalent to the one that would be obtained with more difficulty by bending the fiber to tap some of its signal. The cost-effectiveness of reducing crosstalk must then be compared with that of improving the cladding of the fiber.

Transmission schemes may play a role in anti-jamming and anti-tapping measures. Transmission schemes encompass many techniques, such as modulations that are inured against certain attacks, coding to protect against jamming, intelligent limiting of signals to certain bandwidth and power constraints, and diversity mechanisms which render attacks more difficult. A type of coding scheme which uses diversity is a spread-spectrum technique borrowed from the wireless world. The anti-jam merits of frequency hopping (FH) and direct sequence CDMA (DS-CDMA) [15] for AONs are not simple extensions of the wireless case. The typical wireless jamming problem involves an energy-limited jammer which sends a certain bandwidth over a single link. In AONs, we must contend with jammers which may perform service denial attacks over several links (by using transparency), jam outside the bandwidth over which they transmit (by making use of gain competition in EDFAs) and have few energy limitations. Moreover, CDMA multiple access spread-spectrum techniques are both less bandwidth-efficient and significantly affected by power fluctuations, unless we use more complicated reception and decoding than required for simple orthogonal methods such as WDM or TDM. Since high-speed processing is often expensive and difficult, cooperative multiple access decoding for CDMA may not be as attractive as in other media, and simpler techniques such as TDMA interleaving may be more appropriate.

Architecture and protocols involve such issues as avoidance of easily compromised links for sensitive communications, and judicious wavelength and path assignments which seek to separate trusted users from untrusted users. Consider the following example. If user 1 was observed to burst in power unacceptably, he may be labeled an untrusted user. If network management wants to avoid SD for user 2, it may avoid sharing any OAs or WSSs between users 1 and 2.

In order to be able to separate users, we must have flexibility in our choice of routes and wavelengths. Adding links to a network provides us with more routes over which to route a signal, but may also expose certain parts of the network to new attacks. A common example of optical architecture is the ring. Ring architectures offer two routes between any two points, therefore providing some diversity. A vulnerability of rings is that much of the traffic shares certain links. This may be alleviated somewhat in WDM optical networks. Using WDM over a ring allows different wavelengths to carry traffic between different nodes [16] so that all the traffic does not share the same wavelength, thus reducing certain threats. A mesh architecture gives more flexibility for routing, but also complicates the study of security vulnerabilities.

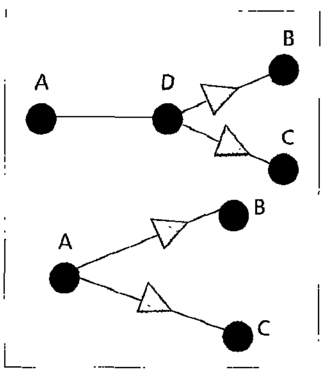

■ Figure 6. *Example of two physical configurations leading to the same logical configuration.*

Detection

Detection of attacks encompasses three functions: event recognition, security failure identification (localization and classification of attack), and alarm generation (appropriate notification to enable adequate reaction to attacks). Note that these three steps need not be consecutive; for instance, alarm generation may be used for localization. Detection of attacks in AONs is more difficult than in traditional electro-optical networks. For some overt attacks, detection may be done independently of the data stream. For instance, rupture of a fiber may be detected if a control signal ceases to be received. Protracted, high-powered jammers may be detected by a power integrator. Covert attacks may be more difficult to detect. Bit error rates (BERs) on fiber are very low (10^{-11} or below); therefore, fairly small interference may cause a relatively large increase in BER, which may not be compensated for by coding schemes designed for lower BERs. A possible countermeasure is to compare two streams that have traversed physically different paths. Such a comparison has the advantage that it may not require knowledge of the particular signaling schemes used, and would detect a problem. Redundant streams would also be useful for reacting to an attack, as described in the next section. However, redundant schemes require more network resources.

Security failure identification of both the location and type of attack may differ significantly from electro-optical networks. Traditional failure identification in networks relies on identifying the domain of an alarm (the components that may lead to an alarm), correlating the alarms, and using algorithms to determine the probability of a certain failure having occurred. Although the same techniques can be applied to AONs, several issues exist [17]. Transparency leads to difficulty in monitoring the nodes. Thus, in [17], only peripheral nodes are tested, and fault propagation is assumed to be instantaneous. If we were to test more nodes in an AON, we may expect that transparency would lead to difficulties in localizing attacks and to many alarms due to the same attack. For instance, a jamming attack could spread transparently and affect several nodes; the point of attack may therefore not be easy to determine. It is not clear whether the delays involved in correlating the alarms and running the algorithms to identify failures may be unacceptably large with respect to the rapid detection necessary, even when these algorithms are polynomial in the size of the network. Preplanned responses to certain sets of alarms may be required.

Alarm generation may fulfill different purposes. The elements in charge of the recovery mechanisms must be notified rapidly. Other nodes in the network may be notified so that they may adaptively change their preplanned response. A signal may be sent to the optical terminal nodes of the AON to warn that the data was possibly corrupted.

Reaction

Reaction consists mainly of isolation of the source of the attack to preclude further attacks, component reconfiguration, rerouting communications, and updating the security status of the network. The reaction to an attack may depend on the location and severity of an attack. For instance, if isolating an attacker cannot be done well because localization is too imprecise or the severity of the attack does not warrant disrupting the network, updating the parameters that determine the level of diversity the network seeks to achieve may be preferable.

As we have mentioned, the extremely high rates associated with AONs may make delays in recovery very critical because large amounts of data may be compromised in a short time. To ensure timely recovery, we may choose to consider preplanned responses to avoid the delays associated with real-time software solutions [18]. A simple preplanned response is the elimination of a user who is seen as a security threat. A user may simply see his communication precluded if the total average power over his assigned wavelength is too high, for instance, by using certain kinds of nonlinear protection switches. Another approach is to consider preplanned routes around compromised network sections. If we determine that a section of fiber or a node is under attack and decide to route around it, we may perform automatic protection switching (APS) to a backup section of the network which is also carrying the required traffic. If the backup traffic is sufficiently delayed with respect to the original traffic, we may not need to retransmit lost or corrupted data. The ability to isolate a fiber to perform APS is afforded by ADMs. An opto-mechanical ADM may switch to a different fiber in a few tens of milliseconds with very low crosstalk. Acousto-optical switches may operate in microseconds and lithium niobate switches in the order of nanoseconds, but such switches exhibit more crosstalk than optomechanical ones. WDM ADMs may be used to add or drop a channel, but also exhibit crosstalk.

The manner in which we perform rerouting will depend on the architecture. Between any two users in a node-redundant network there exist two node-disjoint paths or, in an edge-redundant network, two edge-disjoint paths, and there are several algorithms for finding such routings. WDM allows us to create different logical topologies over a single physical topology. However, analysis of the vulnerability of a logical topology must consider the physical topology and the specific components traversed by a route. For instance, a logical broadcast from a node to two different nodes may be achieved in two different ways [19, 20], but the vulnerabilities of the two physical configurations are different (Fig. 6).

Conclusions

We have given an overview of emerging AONs and their principal components and characteristics. Using this description of AONs, we have presented physical security issues, namely service denial and tapping, in AONs. We have shown how transparency, combined with the characteristics of AON components, creates a set of vulnerabilities which differ significantly from the vulnerabilities found in electronic and electro-optic regenerative networks. While some traditional techniques may be applicable to AONs, countermeasures must be carefully evaluated according to the specific hardware and architecture used. Because service denial attacks are more closely tied to the network physical infrastructure than most of the attacks against which cryptography protects, consideration of service denial attacks prior to widespread deployment of AONs is very important. Understanding the threats could result in small network changes that make the use of both commercial and defense AONs more tenable.

References

[1] H. Onaka *et al.*, "1.1 Tb/s WDM Transmission over a 150 km 1.3 mm Zero-Dispersion Single-Mode Fiber," *Proc. Opt. Fiber Conf. (OFC '96)*, San Jose, CA, Feb. 1996, Postdeadline paper PD19.

[2] A. H. Gnauck *et al.*, "One Terabit/s Transmission Experiment," *Proc. OFC '96*, San Jose, CA, Feb. 1996, Postdeadline paper PD20.

[3] I. P. Kaminow *et al.*, "A Precompetitive Consortium on Wide-band All Optical Networks," *IEEE JSAC*, vol. 14, no. 5, June 1996, pp. 780–99.

[4] C. A Brackett *et al.*, "A Scaleable Multiwavelength Multihop Optical Network: A Proposal for Research on All-optical Networks," *IEEE J. Lightwave Tech.*, vol. 11, no. 5/6, May/June 1993, p. 736.

[5] R. E. Wagner *et al.*, "MONET: Multiwavelength Optical Networking," vol. 14, no. 6, June 1996, pp. 1349–55.

[6] G.-K. Chang *et al.*, "Multiwavelength Reconfigurable WDM/ATM/SONET Network Testbed," *IEEE J. Lightwave Tech.*, vol. 14, no. 6, June 1996, pp. 1320–40.

[7] G. R. Hill *et al.*, "A Transport Network Layer Based on Optical Network Elements," *J. Lightwave Tech.*, vol. 11, no. 5/6, May/June 1993, pp. 667–79.

[8] W. Wells, R. Stone, and E. Miles, "Secure Communications by Optical Homodyne," *IEEE JSAC*, vol. 11, no. 5, June 1993, pp. 770–77.

[9] R. J. Hughes *et al.*, "Quantum Cryptography," *Contemporary Physics*, vol. 36, no. 149, 1995.

[10] S. G. Finn and R. A. Barry, "Optical Services in Future Broadband Networks," *IEEE Network*, Nov./Dec. 1996, vol. 10, no. 6, pp. 7–13.

[11] F. Tosco, *Fiber Optic Communications Handbook*, New York: McGraw Hill, 1990, pp. 237–41.

[12] P. E. Green, *Fiber Optic Networks*, Englewood Cliffs, NJ: Prentice Hall, 1993.

[13] E. Desurvire, *Erbium-Doped Fiber Amplifiers*, New York: John Wiley & Sons, 1994.

[14] A. V. Yakovlev, "An Optical-Fiber System for Transmitting Confidential Information," *Telecommun. and Radio Eng.*, vol. 49, no. 4, 1995.

[15] J. A. Salehi, "Code Division Multiple-Access Techniques in Optical Fiber Networks — Part I: Fundamental Principles," *IEEE Trans. Commun.*, vol. 40, no. 7, July 1992, pp. 1162–70.

[16] A. Hamel *et al.*, "Increased Capacity in an MS Protection Ring Using WDM Technique and OADM: The "Coloured Section" Ring," *Elect. Lett.*, vol. 32, no. 3, 1 Feb. 1996.

[17] I. Katzela, G. Ellinas, and T.E. Stern, "Fault Diagnosis in the Linear Lightwave Network," *1995 Dig. LEOS Summer Topical Meetings*.

[18] T.-H. Wu, *Fiber Network Service Survivability*, New York: Artech House, 1992.

[19] S. Z. Shaikh, "Span-Disjoint Paths for Physical Diversity in Networks," *IEEE Symp. Comp. and Commun.*, 1995, pp. 127–33.

[20] R. Bhandari, "Optimal Diverse Routing in Telecommunication Fiber Networks," *IEEE Infocom 1994*, pp. 11.c.3.1–11.

Biographies

MURIEL MÉDARD is a staff member in the optical communications technology group at MIT. She received her Sc.D. from MIT in 1995. Her thesis was under Prof. R. G. Gallager on the subject of capacity of time-varying channels for multiple access wireless channels. Her M.S. (electrical engineering, MIT, 1991) was done under the supervision of Prof. R. S. Kennedy on scheduling algorithms for optical networks. She received her B.S. in EECS and her B.S. in mathematics in 1989 and her B.S. in humanities in 1991. She received the Vinton-Hayes Fellowship in spring 1991 and the Cronin Fellowship in fall 1991. She worked at NYNEX from 1989 to 1991 as a summer hire and as a consultant examining the performance of multimedia broadband services. She belongs to Tau Beta Pi and Eta Kappa Nu, and is an associate member of Sigma Xi. She is currently working on security in all-optical networks.

DOUGLAS MARQUIS is a staff member in the optical communications technology group at MIT Lincoln Laboratory. He received the B.S. degree in mathematics from the University of Massachusetts in 1982. Since joining Lincoln Laboratory, he has worked on computer simulation of performance for free-space optical communication systems, developed digital control hardware for measurement control, and conducted measurements of optical propagation in the atmosphere. For three years, he worked on the design and implementation of air traffic control automation and radar-based safety systems. He is currently responsible for algorithm design, software development, network management, and applications for the all-optical network program.

RICHARD A. BARRY earned his B.S. (1989), M.S. (1989), and Ph.D. (1993) degrees in electrical engineering form the Laboratory for Information and Decision Systems at MIT. He spent one year (1993-94) on the faculty at The George Washington University, Washington, DC, in the Electrical Engineering and Computer Science Department. Since September1994 he has been a member of the technical staff in the Optical Communications Group, Communications Division, at MIT Lincoln Laboratory. His research interests are in the areas of network architecture and performance, network reliability, information theory, and communication systems. He is a member of Sigma Xi and was awarded the IEEE Communications Society Scholarship in 1990-91. He was also a recipient of the 1993-94 NSF Research Initiation Award. As a student, he worked for Lincoln Laboratory Optical Communications Technology Group and Hughes Aircraft Company.

STEVEN G. FINN [M '69] received his B.S. (1969), M.S. (1969), and ScD. (1975) degrees in electrical engineering from the Massachusetts Institute of Technology. From 1975 to 1980 he worked for Codex Corporation where he held various positions, including director of network processor research and development where he worked on advanced networking products for high-speed data communications networks. Also while at Codex, he was a member of ANSI and CCITT committees involved in public data networking standards development. In 1980, he founded the Bytex Corporation. He held the position of CEO through 1987 and Chairman until he left Bytex in 1990. In 1990, he returned to MIT as a Vinton Hayes Fellow and visiting scientist in the laboratory for information and decision sciences. Currently, he is a principal research scientist and lecturer in the electrical engineering and computer science department at MIT, and is a senior staff member at the MIT Lincoln Laboratory. He is also retained as a consultant for Matrix Partners, a venture capital firm. His current research interests are in the areas of optical networks, high-speed data network transport, network architecture, and network management.